arXiv:1601.03210v1 [cs.CL] 13 Jan 2016

The scarcity of crossing dependencies: a direct outcome of a specific constraint?

Carlos Gómez-Rodríguez
LyS Research Group
Departamento de Computación
Universidade da Coruña
Campus de Elviña, s/n
15071 A Coruña, Spain
carlos.gomez@udc.es

Ramon Ferrer-i-Cancho
Complexity and Quantitative Linguistics Lab
LARCA Research Group
Departament de Ciències de la Computació
Universitat Politècnica de Catalunya (UPC)
Campus Nord, Edifici Omega
08034 Barcelona, Catalonia, Spain
rferrericancho@cs.upc.edu

Abstract

Crossing syntactic dependencies have been observed to be infrequent in natural language, to the point that some syntactic theories and formalisms disregard them entirely. This leads to the question of whether the scarcity of crossings in languages arises from an independent and specific constraint on crossings. We provide statistical evidence suggesting that this is not the case, as the proportion of dependency crossings in a wide range of natural language treebanks can be accurately estimated by a simple predictor based on the local probability that two dependencies cross given their lengths. The relative error of this predictor never exceeds 5% on average, whereas a baseline predictor assuming a random ordering of the words of a sentence incurs a relative error that is at least 6 times greater. Our results suggest that the low frequency of crossings in natural languages is neither originated by hidden knowledge of language nor by the undesirability of crossings *per se*, but as a mere side effect of the principle of dependency length minimization.

1 Introduction

It is well known that **crossing dependencies**, those that cross each other when drawn above the words of a sentence, are relatively uncommon in natural language (Lecerf, 1960; Hays, 1964). Theoretical accounts of languages like English or Japanese often assume that crossing dependencies are nonexistent or negligible, and disallow them in their descriptions of syntax (Sleator and Temperley, 1993; Hudson, 2007; Tanaka, 1997; Kurohashi and Nagao, 1997). In other languages, like Dutch or German, crossing dependencies have been shown to be necessary to describe certain syntactic phenomena (Shieber, 1985; Bresnan et al., 1987), but they still make up only a small percentage of the dependencies found in treebanks (Havelka, 2007; Zeman et al., 2012).

The issue of the presence and frequency of crossing dependencies in the syntax of natural languages has received considerable attention in the computational linguistics community, as supporting them makes parsing computationally harder. Exact inference for parsing models that are restricted to **projective** dependency structures (i.e., those without crossing dependencies or a covered root node[1]) can be solved in polynomial time (Eisner, 1996; Gómez-Rodríguez, Carroll, and Weir, 2008).
In contrast, in the case of general non-projective dependency trees, the problem is only tractable under strong independence assumptions (McDonald and Satta, 2007).

For this reason, many dependency parsers are restricted to projective trees, obtaining

[1]Dependency structures that have no crossing dependencies are called **planar** (Kuhlmann and Nivre, 2006), and are a proper superset of projective structures, as they can include analyses with covered root nodes. However, in practice, structures that are planar but not projective are infrequent in treebanks (Gómez-Rodríguez and Nivre, 2013). Furthermore, if an artificial root node is added at the beginning or end of the sentence, as is common practice in dependency parsing (Ballesteros and Nivre, 2013), then planarity and projectivity become equivalent.

good accuracy in treebanks of languages like English, Chinese or Japanese, where crossing dependencies do not appear or are extremely rare (Chen and Manning, 2014; Iwatate, Asahara, and Matsumoto, 2008). In languages and corpora where they appear more frequently, one option is to resort to approximate inference (Martins, Almeida, and Smith, 2013; Choi and McCallum, 2013). However, the fact that even in such corpora crossing dependencies are still infrequent has enabled an alternative: the so-called mildly non-projective parsers, which can do exact inference supporting very restricted patterns of crossing dependencies that, nevertheless, have been shown to have a high empirical coverage in treebanks (Gómez-Rodríguez, Carroll, and Weir, 2011; Cohen, Gómez-Rodríguez, and Satta, 2011; Pitler, 2014).

In this context, a question naturally arises: what is the reason for the low frequency of crossing dependencies, consistently observed across languages? Answering this question would deepen our understanding of the conditions under which non-projectivity arises, helping us design better formalisms and algorithms to handle it in syntactic parsing and generation. The traditional answer consists of postulating that there is some kind of grammatical ban on non-projectivity (Sleator and Temperley, 1993; Tanaka, 1997; Kurohashi and Nagao, 1997; Starosta, 2003; Lee, 2004; Hudson, 2007). However, this position fails to explain many linguistic phenomena such as topicalization, wh-movement, scrambling or extraposition (Levy et al., 2012; Versley, 2014).

Another option is to assume that crossing dependencies can be grammatical, but only if they follow certain patterns or hard constraints. However, while some mildly non-projective classes of dependency structures have a very good empirical coverage (Kuhlmann and Nivre, 2006; Gómez-Rodríguez and Nivre, 2013), these proposals still face counterexamples that fall outside the restrictions (Chen-Main and Joshi, 2010; Bhat and Sharma, 2012; Chen-Main and Joshi, 2014).

From the perspective of theoretical linguistics, the grammatical ban on non-projectivity can be interpreted

- As a ban set independently from performance considerations, e.g., requiring some hidden parameter to be turned. In this case the ban can be seen as avoidable (e.g., it depends on whether the parameter is on or off for each given language).

- As a consequence of performance constraints associated directly to crossing dependencies. The ban would be inevitable if the cognitive pressures were strong enough but then it would not be properly a ban (a norm added on top of human cognition) but rather a side-effect of cognitive constraints. This view is challenged by psychological and graph theoretic research indicating that crossing dependencies can be easier to process (de Vries et al. (2012) and Ferrer-i-Cancho (2015a) and references therein).

Some researchers have adopted an apparently neutral position concerning the nature of the ban but assume that the low frequency of crossings derives from an independent and specific constraint on crossings: explicitly when postulating a principle of minimization of crossings (Liu, 2008) or implicitly in a large body of research on dependency length minimization that takes for granted that syntactic dependencies should not cross (Liu, 2008; Park and Levy, 2009; Gildea and Temperley, 2010; Futrell, Mahowald, and Gibson, 2015; Gulordava and Merlo, 2015).

If it turned out that non-crossing dependencies can be explained as a side-effect of some cognitive pressure that is not directly associated to crossings (e.g., dependency length minimization), could all these views be regarded as really neutral regarding the nature of the ban?

In this article, we explore a simpler hypothesis: that in order to explain the scarcity of crossing dependencies in language, it is not necessary to assume any underlying rule or principle of human languages that is responsible directly for this fact (including the possibility of some cognitive cost associated directly to crossings). Instead, the low frequency of crossings naturally arises, indirectly, from the actual length of dependencies, which are constrained by a well-known psychological principle: dependency

length minimization (see Ferrer-i-Cancho (2015b) or Tily (2010) for a review).

To this end, we provide statistical evidence that the amount of dependency crossings in a wide range of treebanks can be predicted with small error by a simple estimator based exclusively on dependency length information and information on which edges can potentially cross. The estimator consistently delivers good predictions of the number of crossings, in two different collections of dependency treebanks with diverse annotation criteria. We will argue that this is the best explanation for the low frequency of crossings when both psychological plausibility and parsimony at all levels (from a model of crossings to a general theory of language) are required.

The remainder of the article is organized as follows. Section 2 discusses various ways in which the crossings of a sentence could be predicted. Section 3 presents the predictor of crossings chosen for this article and its theoretical background. The dependency trees used to test the predictor are presented in Section 4. Section 5 shows the results of the predictions, and Section 6 discusses some implications for computational linguistics and linguistic theory.

2 Possible predictors

Here we will examine various possibilities to predict the number of dependency crossings in a sentence. The problem of the origins of non-crossing dependencies can be recast as problem of modeling: we want to find the best model for predicting the number of crossings in a sentence. According to modern model selection, the best model is the one that has the best trade-off between quality of fit (predictive power) and parsimony (Burnham and Anderson, 2002). However, we will be more ambitious, involving further requirements:

- The model must be psychologically realistic. A model that assumes orderings of words that are hard to produce by the human brain should be penalized with respect to one that is based on orderings that real speakers produce (or can rather easily produce). We are not only simply concerned about predicting the low number of crossings of a sentence but also understanding why that number is that low. Hiding the problem under the carpet of grammar or an ad-hoc principle of projectivity does not help.

- Its assumptions must be valid. The predictions of a model may be compatible with real data and even be of high quality but its assumptions may not be supported by real data.

- We are not only concerned about the best model in a local sense but one that leads a general theory of word order or even a comprehensive theory of language that is compact. A real scientific theory is more than a collection of disconnected ideas or models (Bunge, 2001). Models that lead to an unnecessarily fat general theory when integrated into it should also be penalized. Models that exploit assumptions from successful models in other domains should be favored. For instance, a model that allows one to understand not only the scarcity of crossings but also why adjectives tend to be placed after the subject in SOV languages is preferable to one that requires an independent solution to explain the placement of adjectives (Ferrer-i-Cancho, 2015a).

In what follows, we will use C to refer to the number of dependency crossings in the parse of a sentence (i.e., the number of pairs of syntactic dependencies that cross). Our goal is, therefore, to find a suitable predictor for C. For this purpose we will use the notion of the **length** of a dependency, which is defined as the linear distance between the words that are connected: adjacent words are at distance 1, words separated by one word are at distance 2, and so on (Ferrer-i-Cancho, 2004). The sum of the lengths of all dependencies in a sentence will be denoted by D. A **star tree** is a tree with a vertex of maximum degree and $C = 0$ (Ferrer-i-Cancho, 2013b).

2.1 Minimization of crossings

A principle of minimization of crossings (Liu, 2008) leads to a simple deterministic predictor: $C = 0$, reflecting a grammatical ban on non-projectivity (Sleator and Temperley, 1993; Tanaka, 1997; Kurohashi and Nagao, 1997; Starosta, 2003; Lee, 2004; Hudson, 2007).

This predictor is problematic for various reasons:

- Concerning the validity of its assumptions, the model assumes that $C = 0$ independently from D, while C and D are positively correlated in many languages (Ferrer-i-Cancho and Gómez-Rodríguez, 2015).

- Concerning the accuracy of its predictions, this model fails to explain many linguistic phenomena like topicalization, wh-movement, scrambling or extraposition (Levy et al., 2012; Versley, 2014). In many languages, sentences with $C > 0$ are found and then the likelihood of the model is minimum, which indicates that the model is among the worst possible models for crossing dependencies according to modern model selection (Burnham and Anderson, 2002) because its likelihood is zero. Besides, this predictor gives $C = 0$ regardless of the value of D, failing to account for the positive correlation between C and D (Ferrer-i-Cancho and Gómez-Rodríguez, 2015).

- Its psychological validity is unclear. If the model is interpreted as arising from processing difficulties inherent to crossing dependencies (Levy et al., 2012) or computational tractability (as reviewed in Section 1) then it is challenged by psychological and graph theoretic research indicating that sentences with $C > 0$ can be easier to process than sentences with $C = 0$ (see de Vries et al. (2012), Ferrer-i-Cancho (2015a), Chung (1978) and references therein). Another problem is how a language generation process could warrant that $C = 0$. If $C = 0$ is determined before the sentence is produced, how is it possible that sentence production does not introduce (many) crossings? Crossing theory indicates that a star tree is needed to keep a low number of crossings (Ferrer-i-Cancho, 2015a). If $C = 0$ is determined while the sentence is produced (linearized), how are crossings avoided on the fly as real language production is not a batch process (Christiansen and Chater, 2015)? It looks simpler to consider that non-crossing dependencies are a side-effect of a principle of dependency length minimization (Ferrer-i-Cancho, 2006; Ferrer-i-Cancho, 2014a; Ferrer-i-Cancho, 2015a).

- Concerning the compactness of the whole theory, the model $C = 0$ leads to a fatter theory of language because the scarcity of crossings and also the positive correlation between D and C could be explained to a large extent recycling the highly predictive principle of dependency length minimization (Ferrer-i-Cancho, 2015b), as we will see below.

Another option is to assume that crossing dependencies can be grammatical, but only if they follow certain patterns or hard constraints. However, while some mildly non-projective classes of dependency structures have a very good empirical coverage (Kuhlmann and Nivre, 2006; Gómez-Rodríguez and Nivre, 2013), these proposals still face counterexamples that fall outside the restrictions (Chen-Main and Joshi, 2010; Bhat and Sharma, 2012; Chen-Main and Joshi, 2014). One possibility is to relax the simple deterministic predictor above so that on average $C = \alpha(n)$, where α is a constant for sentences of a given length n (Ferrer-i-Cancho and Gómez-Rodríguez, 2015). This allows one to capture the variation in the number of crossings across languages but is problematic for the reasons of the case $\alpha(n) = 0$ that we have examined above. Further arguments can be found in Section 4.3 of Ferrer-i-Cancho (2015a).

2.2 Minimum linear arrangement

A minimum linear arrangement of a sentence is an ordering of the words of the sentence that minimizes the sum of dependency lengths. One may predict the assumed number of crossings by calculating the minimum linear arrangements of a sentence (Ferrer-i-Cancho, 2006). A possible predictor could be the mean number of crossings over all those arrangements.

The predictive power of the model is supported by the fact that solving the minimum linear arrangement problem reduces crossings to practically zero (Ferrer-i-Cancho, 2006), as in many languages. A potential problem of this model is that it has never

been checked whether it predicts the actual number of crossings of real sentences, as far as we know.

Perhaps the major challenge for this predictor is the validity of the assumption of a minimum linear arrangement because:

- The actual value of D in real sentences is located between the minimum and that of a random ordering of vertices (Ferrer-i-Cancho, 2004; Ferrer-i-Cancho and Liu, 2014). The ratio $\Gamma = D/D_{min}$ (where D_{min} is the minimum value of D) is greater than 1.2 in Romanian for sufficiently long sentences (Ferrer-i-Cancho, 2004) and a similar lower bound on language efficiency has been found in English across centuries (Tily, 2010).

- It may not be valid also for theoretical reasons: word order is a multiconstraint satisfaction problem where the principle of dependency length minimization is in conflict with other word order constraints (Ferrer-i-Cancho, 2014b; Futrell, Mahowald, and Gibson, 2015). Thus, a model based on minimum linear arrangements is not that simple: it has to explain why dependency length minimization dominates fully over other principles or provide evidence that the distortion caused by other principles can be neglected. Below we will present a model that does not have this problem because it works on true dependency lengths, which are expected to be determined by the interplay between dependency length minimization and other principles.

- The full minimization of D is cognitively unrealistic, as it is incompatible with the predictions of the now-or-never bottleneck (Christiansen and Chater, 2015). As for the latter, notice that the minimization of D implies that the whole sentence must be available as input for some minimum linear arrangement algorithm, whereas actual language generation and processing is intrinsically online and heavily constrained by our fleeting memory (Christiansen and Chater, 2015).

2.3 Random linear arrangement

If the minimum linear arrangement is too restrictive, one could consider the opposite: predicting the number of crossings assuming a random ordering of the words of the sentence (Ferrer-i-Cancho, 2015a). However, this is cognitively unrealistic: even in languages with high word order flexibility, word order is constrained (Dryer, 2013; Lester and Moscoso del Prado Martín, 2015).
The assumption that the ordering of sentences is arbitrary (unconstrained) is easily rejected by the fact that dependency lengths are below chance in real languages (Ferrer-i-Cancho, 2004; Ferrer-i-Cancho and Liu, 2014;
Futrell, Mahowald, and Gibson, 2015). Furthermore, a random linear arrangement cannot explain the low numbers of crossings observed in real sentences, e.g., a constant low number of crossings requires a star tree (Ferrer-i-Cancho, 2015a). Thus, this predictor is only useful as a random baseline for other predictors. Here we will compare it against a better predictor that is introduced next.

2.4 Random linear arrangement with some knowledge about dependency lengths

A stronger predictor can be built by focusing on the set of pairs of edges that may potentially cross and basing predictions on the actual length of the edges under the assumption of a random arrangement of the four vertices that are potentially involved in an edge crossing (Ferrer-i-Cancho, 2014a). So far, its predictive power is supported by its capacity to predict the actual number of crossings in random trees with an error of about 5% (Ferrer-i-Cancho, 2014a). A crucial goal of the present article is to test the accuracy of its predictions on real sentences.

Concerning assumptions, this model is simpler than the model based on minimum linear arrangements: this model does not assume an unrealistic ordering of the elements of the sentence but the true ordering. This predictor does not assume the principle of dependency length minimization but in case that it predicts the actual number of crossings of real sentences accurately, it success can be easily interpreted as a consequence of the principle of dependency length minimization because actual dependency lengths are below chance

(below $n/3$ approximately), a domain where the probability theory behind this model indicates that a reduction of dependency length yields a reduction in the probability that two edges cross (Ferrer-i-Cancho, 2015a; Ferrer-i-Cancho, 2014a). This predictor is fully compatible with the positive correlation between D and C (Ferrer-i-Cancho and Gómez-Rodríguez, 2015), in contrast with the deterministic predictor ($C = 0$) and its generalization.

Its psychological validity is greater than that of the minimum linear arrangement predictor because it can base its prediction on information from sentences that have actually been produced by a human brain. Contrary to the minimum linear arrangement predictor, this model bases its prediction on true dependency lengths instead of ideal values of D. However, it can be argued that a fundamental assumption of the model, namely that edges are arranged linearly at random, is not supported empirically, for the reasons provided against the random linear arrangement predictor. This is a totally fair criticism, but for this reason this model should be regarded as a null hypothesis rather than as a full realistic model. However, modeling requires a compromise between quality of fit, adequacy and parsimony. If this null hypothesis model provides predictions of sufficient quality on real sentences, do we really need to worry about providing a more realistic but also more complicated model? In the worst case, this predictor is an inevitable baseline for an alternative model.

3 Crossing theory

Here we provide a quick overview of a crossing theory developed in a series of articles (Ferrer-i-Cancho, 2013a; Ferrer-i-Cancho, 2013b; Ferrer-i-Cancho, 2014a; Ferrer-i-Cancho, 2015a). A central concept of crossing theory is Q, the set of pairs of edges of a graph that can potentially cross when their vertices are arranged linearly in some arbitrary order (edges sharing a vertex cannot cross). The cardinality of Q is C_{max}, the **potential number of crossings**. In a tree, one has

$$C_{max} = n(n - 1 - \langle k^2 \rangle)/2, \quad (1)$$

where n is the number of vertices of the tree and $\langle k^2 \rangle$ is the mean of the squared degrees of its vertices (Ferrer-i-Cancho, 2014a). $C_{max} = 0$ if only if the tree is a star tree (Ferrer-i-Cancho, 2013b). C denotes the number of crossings of the linear arrangement of a graph in general while C_{true} denotes the number of crossings of the syntactic dependencies of a real sentence. The **relative number of crossings** is $\bar{C} = C/C_{max}$ or $\bar{C}_{true} = C_{true}/C_{max}$ (Ferrer-i-Cancho, 2014a).

C can be expressed as a sum over Q, i.e.

$$C = \sum_{(e_1,e_2)\in Q} C(e_1, e_2), \quad (2)$$

where $C(e_1, e_2)$ is an indicator variable, $C(e_1, e_2)$ = 1 if the edges e_1 and e_2 cross and $C(e_1, e_2) = 0$ otherwise. The simplest prediction about C than can be made departs from the null hypothesis that the vertices are arranged linearly at random (all possible orderings are equally likely). Following Eq. 2, the expected number of crossings under that null hypothesis turns out to be

$$E_0[C] = \sum_{(e_1,e_2)\in Q} E[C(e_1, e_2)] \quad (3)$$

$$= \sum_{(e_1,e_2)\in Q} p(C(e_1, e_2) = 1), \quad (4)$$

where $p(C(e_1, e_2) = 1)$ is the probability that the edges e_1 and e_2 cross knowing that they belong to Q. Under that null hypothesis, the probability that two edges of Q cross is constant, i.e. $p(C(e_1, e_2) = 1) = 1/3$, yielding (Ferrer-i-Cancho, 2013b) $E_0[C] = C_{max}/3$.

The prediction offered by $E_0[C]$ can be improved by introducing knowledge about the length of the dependencies (edges of length 1 or $n - 1$ are not crossable). The predictor $E_2[C]$ is obtained when $p(C(e_1, e_2) = 1)$ is replaced by $p(C(e_1, e_2) = 1|d(e_1), d(e_2))$, the probability that two arbitrary edges of Q cross, knowing that their lengths are $d(e_1)$ and $d(e_2)$, in Eq. 4, yielding

$$E_2[C] = \sum_{(e_1,e_2)\in Q} p(C(e_1, e_2) = 1|d(e_1), d(e_2)), \quad (5)$$

$p(C(e_1, e_2) = 1|d(e_1), d(e_2))$ depends only on n, $d(e_1)$ and $d(e_2)$ (see Ferrer-i-Cancho (2014a) for further details about $p(C(e_1, e_2) = 1|d(e_1), d(e_2))$). Although $E_0[C]$ and $E_2[C]$ are predictors of C that have the same mathematical structure (they are sums

of probabilities over pairs of edges of Q), $E_0[C]$ is a true expectation while $E_2[C]$ is not.

The **relative error** of a predictor is defined as (Ferrer-i-Cancho, 2014a)

$$\Delta_x = E_x\left[\bar{C}\right] - \bar{C}_{true} = (E_x[C] - C_{true})/C_{max}. \quad (6)$$

Δ_0 will be used as a baseline for Δ_2. Interestingly, Δ_0 converges to $1/3$ for sufficiently long sentences when C_{true} is small (Ferrer-i-Cancho, 2014a).

4 Resources and methods

We considered the corpora in version 2.0 of the HamleDT collection of treebanks (Zeman et al., 2014; Rosa et al., 2014). This collection is a harmonization of existing treebanks for 30 different languages into two well-known annotation styles: Prague dependencies (Hajič et al., 2006) and Universal Stanford dependencies (de Marneffe et al., 2014). Therefore, using this collection allows us to evaluate our predictions of crossings both across a wide range of languages and two popular annotation schemes. The latter is interesting because observations like the number of dependency crossings present in treebank sentences do not only depend on the properties of languages themselves, but also on annotation criteria (Ferrer-i-Cancho and Gómez-Rodríguez (2015) list some examples of how annotation criteria may affect C).

Each of the syntactic dependency structures in the treebanks was preprocessed by removing nodes corresponding to punctuation tokens, as is standard in research related to dependency length (e.g. Ferrer-i-Cancho (2004) and Futrell, Mahowald, and Gibson (2015)), which is only concerned with dependencies between actual words. To preserve the syntactic structure of the rest of the nodes, non-punctuation nodes that had a punctuation node as their head were attached as dependents of their nearest non-punctuation ancestor. Null elements, which appear in the Bengali, Hindi and Telugu corpora, were also subject to the same treatment as punctuation.

After this preprocessing, a syntactic dependency structure was included in our analyses if (1) it defined a tree and (2) the tree was not a star tree. The reason for (1) is that our predictors of the number of crossings assume a tree structure (Ferrer-i-Cancho, 2013b; Ferrer-i-Cancho, 2014a). The reason for (2) is that crossings are impossible in a star tree (Ferrer-i-Cancho, 2013a). Condition (2) implies that the syntactic dependency structure has at least four vertices (otherwise all the possible trees are star trees). By excluding star trees we are discarding trees where the prediction cannot fail. An additional reason for excluding star trees is that their relative number of crossings, C/C_{max}, is not defined because $C = C_{max} = 0$.

Table 1 shows the number of sentences in the original treebanks and the number of sentences actually included in our analyses, after filtering by the criteria (1) and (2) above. The average number of crossings per sentence does not exceed 1 for most of the treebanks.

5 Results

Table 1 shows that the average Δ_2, the relative error of the predictor $E_2[C]$, is small: it does not exceed 5%. Thus, the average Δ_2 is at least 6 times smaller than the baseline $\Delta_0 \approx 30\%$. The averages presented in Table 1 have been produced mixing measurements from sentences of different lengths. This is potentially problematic because the results might be heavily determined by the distribution of sentence lengths (Ferrer-i-Cancho and Liu, 2014). To control for sentence length, sentences were grouped by length and the average Δ_2 was computed for the sentences within each group. Table 2 summarizes the statistical properties over the average Δ_2 of each group. Interestingly, the average over group averages of Δ_2 decreases with respect to the previous analysis: it does not exceed 4.3%. Thus, the average Δ_2 is at least 7 times smaller than the baseline error, again $\Delta_0 \approx 30\%$. The minimum size of a group is one sentence; the qualitative results are very similar if the minimum size is set to 2.

6 Discussion

We have shown that $E_2[C]$ predicts C_{true} with small error, much better than the baseline. The positive results are not surprising given the previous success of $E_2[C]$ predicting crossings on uniformly random trees, where Δ_2 is about 5%, i.e. about 6 times smaller than the baseline Δ_0, for sufficiently

Table 1: Summary of results for each treebank: number of sentences before and after filtering, average values of C_{true} and Δ_0, and average, median and standard deviation of the relative error Δ_2, over the trees of each treebank. Romanian (Prague) is the only treebank with no crossing dependencies.

Treebank	#Sent	#Sent (filtered)	C_{true} (avg.)	Δ_0 (avg.)	Δ_2 (avg.)	Δ_2 (median)	Δ_2 (st. dev.)
Stanford annotation							
Arabic	7547	2280	0.9807	0.328	0.019	0.016	0.021
Basque	11225	9072	0.1391	0.330	0.028	0.022	0.033
Bengali	1129	678	0.1062	0.321	0.027	0.000	0.051
Bulgarian	13221	12119	0.3598	0.326	0.045	0.039	0.042
Catalan	14924	14520	0.6419	0.331	0.034	0.029	0.023
Czech	87913	74843	0.5277	0.326	0.040	0.035	0.035
Danish	5512	4894	0.6800	0.324	0.047	0.040	0.038
Dutch	13735	10974	1.3980	0.311	0.046	0.041	0.051
English	18791	18275	0.5241	0.330	0.049	0.043	0.031
Estonian	1315	851	0.0376	0.331	0.016	0.000	0.037
Finnish	4307	4078	0.3183	0.326	0.034	0.028	0.038
German	38020	33492	0.7826	0.325	0.050	0.046	0.036
Greek	2902	2584	0.6540	0.330	0.039	0.033	0.028
Anc. Greek	21173	18713	3.2621	0.244	0.030	0.027	0.058
Hindi	13274	12417	0.3043	0.332	0.027	0.025	0.017
Hungarian	6424	6103	0.9720	0.326	0.036	0.031	0.033
Italian	3359	2502	0.4153	0.329	0.035	0.029	0.032
Japanese	17753	4614	0.1641	0.326	0.024	0.019	0.032
Latin	3473	3036	2.1785	0.282	0.034	0.031	0.046
Persian	12455	11579	0.5914	0.326	0.027	0.023	0.031
Portuguese	9359	8621	0.6336	0.328	0.039	0.033	0.032
Romanian	4042	3145	0.1021	0.331	0.028	0.021	0.036
Russian	34895	31581	0.4171	0.326	0.038	0.032	0.035
Slovak	57408	47727	0.4559	0.324	0.044	0.036	0.044
Slovenian	1936	1719	0.7749	0.322	0.047	0.039	0.046
Spanish	15984	15354	0.6218	0.331	0.034	0.029	0.024
Swedish	11431	10714	0.4871	0.328	0.043	0.039	0.034
Tamil	600	584	0.0240	0.333	0.026	0.022	0.025
Telugu	1450	429	0.0140	0.322	0.016	0.000	0.045
Turkish	5935	3862	0.0984	0.330	0.031	0.025	0.038
Prague annotation							
Arabic	7547	2248	0.0881	0.333	0.013	0.010	0.016
Basque	11225	8717	0.1252	0.330	0.026	0.021	0.029
Bengali	1129	651	0.1244	0.320	0.025	0.000	0.052
Bulgarian	13221	11947	0.1248	0.329	0.023	0.017	0.029
Catalan	14924	14556	0.0873	0.333	0.020	0.017	0.016
Czech	87913	70023	0.3729	0.327	0.031	0.025	0.031
Danish	5512	4840	0.1643	0.331	0.027	0.022	0.027
Dutch	13735	11131	0.9898	0.315	0.034	0.027	0.041
English	18791	18369	0.1072	0.333	0.034	0.029	0.024
Estonian	1315	843	0.0130	0.332	0.013	0.000	0.031
Finnish	4307	4011	0.1279	0.330	0.028	0.024	0.031
German	38020	32443	0.7230	0.326	0.043	0.039	0.033
Greek	2902	2543	0.2057	0.332	0.030	0.024	0.023
Anc. Greek	21173	16237	3.3528	0.243	0.025	0.020	0.058
Hindi	13274	12334	0.3875	0.330	0.015	0.012	0.015
Hungarian	6424	5047	0.8675	0.326	0.034	0.030	0.032
Italian	3359	2398	0.0621	0.333	0.020	0.014	0.024
Japanese	17753	4792	0.0002	0.333	0.006	0.000	0.013
Latin	3473	2833	1.8503	0.286	0.036	0.032	0.047
Persian	12455	11632	0.4024	0.329	0.030	0.024	0.033
Portuguese	9359	8596	0.2465	0.331	0.021	0.016	0.021
Romanian	4042	3193	0.0000	0.333	0.015	0.005	0.026
Russian	34895	31900	0.1570	0.330	0.027	0.021	0.028
Slovak	57408	44297	0.2688	0.326	0.034	0.026	0.039
Slovenian	1936	1581	0.3125	0.327	0.035	0.027	0.038
Spanish	15984	15424	0.1105	0.333	0.020	0.016	0.017
Swedish	11431	10207	0.1946	0.332	0.034	0.029	0.029
Tamil	600	585	0.0137	0.333	0.023	0.019	0.023
Telugu	1450	373	0.0080	0.325	0.014	0.000	0.043
Turkish	5935	3518	0.1373	0.327	0.015	0.000	0.026

Table 2: Summary of results for each treebank: number of distinct sentence lengths, average Δ_0 and, average, median and standard deviation of the average values of Δ_2 over the groups of sentences with the same length.

Treebank	#Lengths	Δ_0 (avg.)	Δ_2 (avg.)	Δ_2 (median)	Δ_2 (st. dev.)
Stanford annotation					
Arabic	109	0.331	0.014	0.013	0.006
Basque	35	0.331	0.026	0.022	0.017
Bengali	18	0.322	0.034	0.028	0.026
Bulgarian	64	0.330	0.029	0.027	0.012
Catalan	98	0.332	0.023	0.021	0.008
Czech	88	0.330	0.024	0.022	0.010
Danish	66	0.328	0.031	0.029	0.013
Dutch	54	0.319	0.037	0.035	0.019
English	74	0.331	0.033	0.031	0.013
Estonian	25	0.331	0.036	0.032	0.020
Finnish	41	0.329	0.028	0.025	0.016
German	85	0.328	0.033	0.032	0.012
Greek	75	0.331	0.027	0.024	0.010
Anc. Greek	66	0.293	0.025	0.024	0.019
Hindi	69	0.332	0.020	0.018	0.008
Hungarian	65	0.328	0.027	0.025	0.014
Italian	69	0.331	0.024	0.022	0.010
Japanese	44	0.330	0.021	0.020	0.010
Latin	59	0.309	0.031	0.029	0.018
Persian	93	0.329	0.023	0.021	0.009
Portuguese	88	0.331	0.024	0.023	0.009
Romanian	46	0.332	0.023	0.021	0.011
Russian	80	0.330	0.024	0.022	0.010
Slovak	92	0.330	0.024	0.022	0.010
Slovenian	57	0.326	0.034	0.032	0.016
Spanish	95	0.332	0.023	0.022	0.009
Swedish	74	0.329	0.028	0.026	0.011
Tamil	40	0.333	0.023	0.020	0.011
Telugu	10	0.330	0.043	0.030	0.030
Turkish	51	0.332	0.030	0.027	0.013
Prague annotation					
Arabic	109	0.333	0.010	0.008	0.005
Basque	35	0.331	0.024	0.021	0.015
Bengali	17	0.321	0.034	0.027	0.030
Bulgarian	63	0.332	0.016	0.014	0.009
Catalan	98	0.333	0.014	0.012	0.006
Czech	87	0.331	0.019	0.017	0.009
Danish	66	0.332	0.019	0.017	0.010
Dutch	54	0.323	0.027	0.024	0.016
English	75	0.333	0.023	0.021	0.010
Estonian	25	0.332	0.032	0.028	0.019
Finnish	41	0.331	0.024	0.021	0.013
German	85	0.329	0.029	0.027	0.011
Greek	74	0.333	0.021	0.019	0.008
Anc. Greek	65	0.292	0.021	0.021	0.020
Hindi	69	0.331	0.012	0.010	0.007
Hungarian	65	0.329	0.026	0.023	0.013
Italian	68	0.333	0.014	0.012	0.008
Japanese	44	0.333	0.008	0.006	0.007
Latin	59	0.313	0.031	0.030	0.017
Persian	93	0.331	0.022	0.021	0.009
Portuguese	88	0.332	0.013	0.012	0.006
Romanian	46	0.333	0.012	0.010	0.008
Russian	80	0.332	0.017	0.016	0.008
Slovak	87	0.331	0.020	0.018	0.010
Slovenian	50	0.329	0.027	0.024	0.014
Spanish	95	0.333	0.014	0.013	0.006
Swedish	73	0.331	0.022	0.021	0.010
Tamil	40	0.333	0.018	0.016	0.010
Telugu	10	0.331	0.037	0.029	0.033
Turkish	49	0.331	0.015	0.013	0.010

long sentences (Ferrer-i-Cancho, 2014a). It is also worth noting that $E_2[C]$ behaves well even in the treebanks with the lowest proportion of crossings, where one could argue that grammar would impose the heaviest constraints against crossings. For example, it achieves a particularly low relative error in the Romanian and Japanese Prague treebanks although they contain no or almost no crossings (Table 1).

It could be argued that the good predictions of $E_2[C]$ are not surprising at all because the syntactic dependency structures that we have analyzed could be the result of some sophisticated apparatus: a complex language faculty or external grammatical knowledge which could have produced, indirectly, a distribution of dependency lengths and vertex degrees that is favorable for $E_2[C]$. Then the input with which the predictor yields good predictions, e.g., dependency lengths, would be an indirect result of that complex device. However, $E_2[C]$ does not require such a device: $E_2[C]$ also makes accurate predictions on uniformly random trees with a small number of crossings (Ferrer-i-Cancho, 2014a). Therefore, the need of external grammatical knowledge to explain the origins of non-crossing dependencies is seriously challenged.

The high precision of $E_2[C]$ suggests that the actual number of crossings in sentences might be a side effect of the dependency lengths, which are in turn constrained by a general principle of dependency length minimization (see Ferrer-i-Cancho (2015b) for a review of the empirical and theoretical backup of that principle). A ban of crossings by grammar (e.g., Hudson (2007; Tanaka (1997)), a principle of minimization of crossings (Liu, 2008) or a competence-plus (Hurford, 2012) limiting the number of crossings, may not be necessary to explain the low frequency of crossings in world languages.

In spite of the arguments in favor of a model predicting crossings based on dependency lengths reviewed and expanded in this article, it looks difficult to rule out some principle of minimization of crossings or projectivity constraint. The reason is the positive correlation between crossings and dependency lengths that has been unveiled by this article and previous research combining both theory and experiment (see Ferrer-i-Cancho (2014a), Ferrer-i-Cancho (2015a), Ferrer-i-Cancho and Gómez-Rodríguez (2015) and references therein). The question is: what is the causal force for the scarcity of crossings: (a) a principle of minimization of crossings that explains why dependency lengths are short or (b) a principle of dependency length minimization that explains the scarcity of crossings? (Ferrer-i-Cancho, 2015a). A temporary solution to this dilemma is straightforward if we are seriously concerned about the construction of a general theory of language that is not only highly predictive but also parsimonious: a theory of language based on (b) is more parsimonious than one based on (a) (Ferrer-i-Cancho, 2015a).

Our results also have interesting implications for the field of parsing, and related tasks like machine translation (Gómez-Rodríguez et al., 2009), where algorithms and formalisms face a trade-off between coverage of crossing dependencies and computational cost. If the low frequency of crossing dependencies arises naturally from dependency length minimization, this suggests that the adequacy of the so-called mildly non-projective parsers and formalisms, which support restricted forms of non-projectivity, should be evaluated solely on their empirical coverage in natural language treebanks in relation to their computational efficiency, rather than the conformance to some concrete set of grammatical rules or principles governing which kinds of crossing dependencies should be allowed.

This is relevant because the latter has typically been taken into account in the literature: some examples are the limitations like constant growth in the classic characterization of mild context-sensitivity (Joshi, 1985; Kallmeyer, 2010); or the popularity of well-nested dependency structures with gap degree at most 1 (Kuhlmann and Nivre, 2006), probably the best-known class of mildly non-projective structures (see for example Pitler, Kannan, and Marcus (2012)) over alternatives like mildly ill-nested dependency structures of gap degree at most 1, which have a larger coverage with the exact same parsing complexity (Gómez-Rodríguez, Carroll, and Weir, 2011). On the other hand, as an example of the purely empirical approach, Gildea (2010) shows how the optimal factorization of LCFRS rules in terms of time complexity does not necessarily arise from observing restrictions like well-nestedness or even

minimum fan-out (gap degree).

The good empirical coverage of mildly non-projective classes of structures defined by such diverse restrictions as well-nestedness (Kuhlmann and Nivre, 2006), 1-endpoint-crossing (Pitler, Kannan, and Marcus, 2013; Pitler, 2014), 2-planarity (Gómez-Rodríguez and Nivre, 2013), parsability with the basic Attardi transitions (Attardi, 2006; Cohen, Gómez-Rodríguez, and Satta, 2011) or even more so mildly ill-nestedness (Gómez-Rodríguez, Carroll, and Weir, 2011) or 1-ill-nestedness (Maier and Lichte, 2011) may be explained simply because of their flexibility, not because the particular restrictions on crossings they introduce necessarily match real underlying restrictions in a direct fashion. Our findings suggest that realistic constraints on dependency lengths could help to define, maybe probabilistically, the mild non-projectivity that is needed to describe and parse real languages with an affordable cost.

7 Acknowledgments

RFC is funded by the grants 2014SGR 890 (MACDA) from AGAUR (Generalitat de Catalunya) and also the APCOM project (TIN2014-57226-P) from MINECO (Ministerio de Economia y Competitividad). CGR is partially funded by the TELEPARES-UDC project (FFI2014-51978-C2-2-R) from MINECO and an Oportunius program grant from Xunta de Galicia.

We thank Morten Christiansen for helpful discussions, Wolfgang Maier for comments on an earlier version of this manuscript, and Dan Zeman for help with data conversion.

References

[Attardi2006] Attardi, Giuseppe. 2006. Experiments with a multilanguage non-projective dependency parser. In *Proceedings of the 10th Conference on Computational Natural Language Learning (CoNLL)*, pages 166–170.

[Ballesteros and Nivre2013] Ballesteros, Miguel and Joakim Nivre. 2013. Going to the roots of dependency parsing. *Computational Linguistics*, 39(1):5–13.

[Bhat and Sharma2012] Bhat, Riyaz Ahmad and Dipti Misra Sharma. 2012. Non-projective structures in Indian language treebanks. In *Proceedings of the Eleventh International Workshop on Treebanks and Linguistic Theories (TLT11)*, pages 25–30, Lisbon, Portugal. Edições Colibri.

[Bresnan et al.1987] Bresnan, Joan, Ronald M. Kaplan, Stanley Peters, and Annie Zaenen. 1987. Cross-serial dependencies in Dutch. In W. J. Savitch, E. Bach, W. Marsh, and G. Safran-Naveh, editors, *The Formal Complexity of Natural Language*. Reidel, Dordrecht, pages 286–319.

[Bunge2001] Bunge, Mario. 2001. *La science, sa méthode et sa philosophie*. Vigdor, Paris.

[Burnham and Anderson2002] Burnham, Kenneth P. and David R. Anderson. 2002. *Model selection and multimodel inference. A practical information-theoretic approach*. Springer, New York, 2nd edition.

[Chen and Manning2014] Chen, Danqi and Christopher Manning. 2014. A fast and accurate dependency parser using neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 740–750, Doha, Qatar.

[Chen-Main and Joshi2010] Chen-Main, Joan and Aravind K Joshi. 2010. Unavoidable ill-nestedness in natural language and the adequacy of tree local-MCTAG induced dependency structures. In *Proceedings of the Tenth International Workshop on Tree Adjoining Grammar and Related Formalisms (TAG+ 10)*, pages 119–126.

[Chen-Main and Joshi2014] Chen-Main, Joan and Aravind K. Joshi. 2014. A dependency perspective on the adequacy of tree local multi-component tree adjoining grammar. *J. Log. Comput.*, 24(5):989–1022.

[Choi and McCallum2013] Choi, Jinho D. and Andrew McCallum. 2013. Transition-based dependency parsing with selectional branching. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1052–1062, Sofia, Bulgaria.

[Christiansen and Chater2015] Christiansen, Morten H. and Nick Chater. 2015. The now-or-never bottleneck: a fundamental constraint on language. *Behavioral and Brain Sciences*, page in press.

[Chung1978] Chung, Fan R. K. 1978. A conjectured minimum valuation tree. *SIAM Review*, 20:601 604.

[Cohen, Gómez-Rodríguez, and Satta2011] Cohen, Shay B., Carlos Gómez-Rodríguez, and Giorgio Satta. 2011. Exact inference for generative probabilistic non-projective dependency parsing. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1234–1245, Edinburgh, Scotland, UK.

[de Marneffe et al.2014] de Marneffe, Marie-Catherine, Timothy Dozat, Natalia Silveira, Katri Haverinen,

Filip Ginter, Joakim Nivre, and Christopher D. Manning. 2014. Universal Stanford dependencies: a cross-linguistic typology. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Hrafn Loftsson, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland, may. European Language Resources Association (ELRA).

[de Vries et al.2012] de Vries, Meinou H., Karl M. Petersson, Sebastian Geukes, Pienie Zwitserlood, and Morten H. Christiansen. 2012. Processing multiple non-adjacent dependencies: evidence from sequence learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1598):2065–2076.

[Dryer2013] Dryer, Matthew S. 2013. Order of subject, object and verb. In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.

[Eisner1996] Eisner, Jason. 1996. Three new probabilistic models for dependency parsing: An exploration. In *Proceedings of the 16th International Conference on Computational Linguistics (COLING-96)*, pages 340–345, San Francisco, CA, USA.

[Ferrer-i-Cancho2004] Ferrer-i-Cancho, Ramon. 2004. Euclidean distance between syntactically linked words. *Physical Review E*, 70:056135.

[Ferrer-i-Cancho2006] Ferrer-i-Cancho, Ramon. 2006. Why do syntactic links not cross? *Europhysics Letters*, 76(6):1228–1235.

[Ferrer-i-Cancho2013a] Ferrer-i-Cancho, Ramon. 2013a. Hubiness, length, crossings and their relationships in dependency trees. *Glottometrics*, 25:1–21.

[Ferrer-i-Cancho2013b] Ferrer-i-Cancho, Ramon. 2013b. Random crossings in dependency trees. *http://arxiv.org/abs/1305.4561*.

[Ferrer-i-Cancho2014a] Ferrer-i-Cancho, Ramon. 2014a. A stronger null hypothesis for crossing dependencies. *Europhysics Letters*, 108:58003.

[Ferrer-i-Cancho2014b] Ferrer-i-Cancho, Ramon. 2014b. Why might SOV be initially preferred and then lost or recovered? A theoretical framework. In E. A. Cartmill, S. Roberts, H. Lyn, and H. Cornish, editors, *THE EVOLUTION OF LANGUAGE - Proceedings of the 10th International Conference (EVOLANG10)*, pages 66–73, Vienna, Austria. Wiley. Evolution of Language Conference (Evolang 2014), April 14-17.

[Ferrer-i-Cancho2015a] Ferrer-i-Cancho, Ramon. 2015a. Non-crossing dependencies: least effort, not grammar. In A. Mehler, A. Lücking, S. Banisch, P. Blanchard, and B. Job, editors, *Towards a theoretical framework for analyzing complex linguistic networks*. Springer, Berlin, pages 203–234.

[Ferrer-i-Cancho2015b] Ferrer-i-Cancho, Ramon. 2015b. The placement of the head that minimizes online memory. A complex systems approach. *Language Dynamics and Change*, 5:141–164.

[Ferrer-i-Cancho and Gómez-Rodríguez2015] Ferrer-i-Cancho, Ramon and Carlos Gómez-Rodríguez. 2015. Crossings as a side effect of dependency lengths. *http://arxiv.org/abs/1508.06451*.

[Ferrer-i-Cancho and Liu2014] Ferrer-i-Cancho, Ramon and Haitao Liu. 2014. The risks of mixing dependency lengths from sequences of different length. *Glottotheory*, 5:143–155.

[Futrell, Mahowald, and Gibson2015] Futrell, Richard, Kyle Mahowald, and Edward Gibson. 2015. Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences*, 112(33):10336–10341.

[Gildea2010] Gildea, Daniel. 2010. Optimal parsing strategies for linear context-free rewriting systems. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, HLT '10, pages 769–776, Stroudsburg, PA, USA. Association for Computational Linguistics.

[Gildea and Temperley2010] Gildea, Daniel and David Temperley. 2010. Do grammars minimize dependency length? *Cognitive Science*, 34(2):286–310.

[Gómez-Rodríguez, Carroll, and Weir2008] Gómez-Rodríguez, Carlos, John Carroll, and David Weir. 2008. A deductive approach to dependency parsing. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL'08:HLT)*, pages 968–976.

[Gómez-Rodríguez, Carroll, and Weir2011] Gómez-Rodríguez, Carlos, John Carroll, and David Weir. 2011. Dependency parsing schemata and mildly non-projective dependency parsing. *Computational Linguistics*, 37(3):541–586.

[Gómez-Rodríguez et al.2009] Gómez-Rodríguez, Carlos, Marco Kuhlmann, Giorgio Satta, and David J. Weir. 2009. Optimal reduction of rule length in linear context-free rewriting systems. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 539–547, Boulder, CO, USA.

[Gómez-Rodríguez and Nivre2013] Gómez-Rodríguez, Carlos and Joakim Nivre. 2013. Divisible transition systems and multiplanar dependency parsing. *Computational Linguistics*, 39(4):799–845.

[Gulordava and Merlo2015] Gulordava, Kristina and Paola Merlo. 2015. Diachronic trends in word order freedom and dependency length in dependency-annotated corpora of Latin and ancient Greek. In *Proceedings of the Third International Conference on Dependency Linguistics (Depling 2015)*, pages 121–130, Uppsala, Sweden. Uppsala University.

[Hajič et al.2006] Hajič, Jan, Jarmila Panevová, Eva Hajičová, Jarmila Panevová, Petr Sgall, Petr Pajas, Jan Štěpánek, Jiří Havelka, and Marie Mikulová. 2006. Prague Dependency Treebank 2.0. CDROM CAT: LDC2006T01, ISBN 1-58563-370-4. Linguistic Data Consortium.

[Havelka2007] Havelka, Jiří. 2007. Beyond projectivity: Multilingual evaluation of constraints and measures on non-projective structures. In *ACL 2007: Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics*, pages 608–615.

[Hays1964] Hays, David G. 1964. Dependency theory: A formalism and some observations. *Language*, 40:511–525.

[Hudson2007] Hudson, Richard A. 2007. *Language Networks: The New Word Grammar*. Oxford University Press.

[Hurford2012] Hurford, James R., 2012. *Chapter 3. Syntax in the Light of Evolution*, pages 175–258. Oxford University Press, Oxford.

[Iwatate, Asahara, and Matsumoto2008] Iwatate, Masakazu, Masayuki Asahara, and Yuji Matsumoto. 2008. Japanese dependency parsing using a tournament model. In *Proceedings of the 22Nd International Conference on Computational Linguistics - Volume 1*, pages 361–368.

[Joshi1985] Joshi, Aravind. 1985. How much context-sensisitivity is necessary for characterizing structural descriptions – tree adjoining grammars. In David Dowty, Lauri Karttunen, and Arnold Zwicky, editors, *Natural Language Processing: Psycholinguistic, Computational and Theoretical Perspectives*. Cambridge University Press, pages 206–250.

[Kallmeyer2010] Kallmeyer, Laura. 2010. *Parsing Beyond Context-Free Grammars*. Springer Publishing Company, Incorporated, 1st edition.

[Kuhlmann and Nivre2006] Kuhlmann, Marco and Joakim Nivre. 2006. Mildly non-projective dependency structures. In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 507–514.

[Kurohashi and Nagao1997] Kurohashi, Sadao and Makoto Nagao. 1997. Kyoto university text corpus project (in Japanese). In *Proceedings of the 3rd Annual Meeting of the Association for Natural Language Processing*, pages 115–118.

[Lecerf1960] Lecerf, Yves. 1960. Programme des conflits, modèle des conflits. *Bulletin bimestriel de l'ATALA*, 1:(4): 11–18, (5): 17–36.

[Lee2004] Lee, Songwook. 2004. A statistical model for identifying grammatical relations in korean sentences. *IEICE Transactions*, 87-D(12):2863–2871.

[Lester and Moscoso del Prado Martín2015] Lester, Nicolas A. and Fermín Moscoso del Prado Martín. 2015. Word order in a grammarless language: A 'small-data' information-theoretic approach. In D.C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, and P. P. Maglio, editors, *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*, pages 1314–1319, Austin, TX. Cognitive Science Society.

[Levy et al.2012] Levy, Roger, Evelina Fedorenko, Mara Breen, and Ted Gibson. 2012. The processing of extraposed structures in English. *Cognition*, 122(1):12 – 36.

[Liu2008] Liu, Haitao. 2008. Dependency distance as a metric of language comprehension difficulty. *Journal of Cognitive Science*, 9:159–191.

[Maier and Lichte2011] Maier, Wolfgang and Timm Lichte. 2011. Characterizing discontinuity in constituent treebanks. In *Formal Grammar. 14th International Conference, FG 2009. Bordeaux, France, July 25-26, 2009. Revised Selected Papers*, volume 5591 of *LNCS/LNAI*, pages 167–182, Berlin-Heidelberg. Springer.

[Martins, Almeida, and Smith2013] Martins, Andre, Miguel Almeida, and Noah A. Smith. 2013. Turning on the turbo: Fast third-order non-projective turbo parsers. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 617–622, Sofia, Bulgaria.

[McDonald and Satta2007] McDonald, Ryan and Giorgio Satta. 2007. On the complexity of non-projective data-driven dependency parsing. In *IWPT 2007: Proceedings of the 10th International Conference on Parsing Technologies*, pages 121–132.

[Park and Levy2009] Park, Y. Albert and Roger Levy. 2009. Minimal-length linearizations for mildly context-sensitive dependency trees. In *Proceedings of the 10th Annual Meeting of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT) conference*, pages 335–343, Stroudsburg, PA, USA. Association for Computational Linguistics.

[Pitler2014] Pitler, Emily. 2014. A crossing-sensitive third-order factorization for dependency parsing. *Transactions of the Association for Computational Linguistics*, 2:41–54.

[Pitler, Kannan, and Marcus2012] Pitler, Emily, Sampath Kannan, and Mitchell Marcus. 2012. Dynamic programming for higher order parsing of gap-minding trees. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 478–488, Jeju Island, Korea, July. Association for Computational Linguistics.

[Pitler, Kannan, and Marcus2013] Pitler, Emily, Sampath Kannan, and Mitchell Marcus. 2013. Finding optimal 1-endpoint-crossing trees. *Transactions of the Association of Computational Linguistics*, 1:13–24.

[Rosa et al.2014] Rosa, Rudolf, Jan Maek, David Marecek, Martin Popel, Daniel Zeman, and Zdenek abokrtsk. 2014. HamleDT 2.0: Thirty dependency treebanks stanfordized. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Hrafn Loftsson, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland, may. European Language Resources Association (ELRA).

[Shieber1985] Shieber, Stuart M. 1985. Evidence against the context-freeness of natural language. *Linguistics and Philosophy*, 8(3):333–343.

[Sleator and Temperley1993] Sleator, Daniel and Davy Temperley. 1993. Parsing English with a link grammar. In *Proceedings of the Third International Workshop on Parsing Technologies (IWPT)*, pages 277–292.

[Starosta2003] Starosta, Stanley. 2003. Lexicase grammar. In Vilmos Ágel, Ludwig M. Eichinger, Hans Werner Eroms, Peter Hellwig, Hans Jürgen Heringer, and Henning Lobin, editors, *Dependency and Valency – An International Handbook of Contemporary Research*. Walter de Gruyter, pages 526–545.

[Tanaka1997] Tanaka, Hidekazu. 1997. Invisible movement in sika-nai and the linear crossing constraint. *Journal of East Asian Linguistics*, 6:143–188.

[Tily2010] Tily, Harry J. 2010. *The role of processing complexity in word order variation and change*. Ph.D. thesis, Stanford University, August. Chapter 3: Dependency lengths.

[Versley2014] Versley, Yannick. 2014. Experiments with easy-first nonprojective constituent parsing. In *Proceedings of the First Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical Languages*, pages 39–53, Dublin, Ireland, August. Dublin City University.

[Zeman et al.2014] Zeman, Daniel, Ondřej Dušek, David Mareček, Martin Popel, Loganathan Ramasamy, Jan Štěpánek, Zdeněk Žabokrtský, and Jan Hajič. 2014. HamleDT: Harmonized multi-language dependency treebank. *Language Resources and Evaluation*, 48(4):601–637.

[Zeman et al.2012] Zeman, Daniel, David Marecek, Martin Popel, Loganathan Ramasamy, Jan Łtepnek, Zdenek abokrtsk, and Jan Hajic. 2012. HamleDT: To parse or not to parse? In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey. ELRA.